

Near-Optimal Multi-Unit Auctions with Ordered Bidders

Elias Koutsoupias*
University of Oxford and
University of Athens
elias@cs.ox.ac.uk

Stefano Leonardi†
Sapienza University of Rome
leon@dis.uniroma1.it

Tim Roughgarden‡
Stanford University
tim@cs.stanford.edu

December 13, 2012

Abstract

We construct prior-free auctions with constant-factor approximation guarantees with ordered bidders, in both unlimited and limited supply settings. We compare the expected revenue of our auctions on a bid vector to the monotone price benchmark, the maximum revenue that can be obtained from a bid vector using supply-respecting prices that are nonincreasing in the bidder ordering and bounded above by the second-highest bid. As a consequence, our auctions are simultaneously near-optimal in a wide range of Bayesian multi-unit environments.

1 Introduction

The goal in prior-free auction design is to design auctions that have robust, input-by-input performance guarantees. Traditionally, auctions are evaluated using average-case or Bayesian analysis, and expected auction

*Partially supported by the European Union Seventh Framework Programme FP7, grant 284731 (UaESMC), and by the ESF-NSRF research program Thales (AGT).

†This work was partially supported from EU ERC research grant PAAI (Practical Approximation Algorithms)

‡Supported in part by NSF grant CCF-1016885, an ONR PECASE Award, and an AFOSR MURI grant.

performance is optimized with respect to a prior distribution over inputs (i.e., bid vectors). The Bayesian versions of the problems we consider are completely solved [19]. Worst-case guarantees are desirable when, for example, good prior information is expensive or impossible to acquire, and when a single auction is to be re-used several times, in settings with different or not-yet-known input distributions.

Prior-free auctions were first studied by Goldberg et al. [9, 10]. They focused on symmetric settings, where goods and bidders are identical, and sought auctions with expected revenue close to the fixed-price benchmark $\mathcal{F}^{(2)}$, defined as the maximum revenue that can be obtained from a given bid vector by offering every bidder a common posted price (i.e., take-it-or-leave-it offer) that is at most the second-highest bid. Goldberg et al. [9] showed that no auction has expected revenue more than a $\approx .42$ fraction of $\mathcal{F}^{(2)}$ for every bid vector, and constructed auctions with expected revenue at least a constant fraction of this benchmark on every input. See Hartline and Karlin [12] for a survey of further work in this vein.

Hartline and Roughgarden [13] proposed a framework for defining meaningful performance benchmarks much more generally — when bidders or feasibility constraints are asymmetric, and for objective functions other than revenue. The first step of this framework is a “Bayesian thought experiment” — if bidders’ valuations were drawn from a prior distribution in some class, what would the optimal auction be? The second step is to characterize the collection \mathcal{C} of all optimal auctions that can arise, ranging over all permissible prior distributions. Finally, given a bid vector \mathbf{b} , the performance benchmark is defined as the maximum objective function value obtained by an auction in \mathcal{C} on the input \mathbf{b} . This framework regenerates the $\mathcal{F}^{(2)}$ benchmark (modulo the technically necessary upper bound on prices) and has been used for several other objective functions and asymmetric environments [5, 13, 14, 15, 17]. Every benchmark generated by this framework is automatically well motivated in the following sense: if the performance of an auction is within a constant factor of such a benchmark for every input, then in particular it is simultaneously near-optimal in every Bayesian environment with valuations drawn from one of the permissible prior distributions.¹

Leonardi and Roughgarden [17] studied the design and analysis of prior-free digital goods (i.e., unlimited supply) auctions with asymmetric bid-

¹This weaker goal of good *prior-independent* auctions can also be studied in its own right [6, 7, 21]. See [3, 4, 18] for other interpolations between average-case and worst-case analysis of auctions.

ders. They pointed out that the framework in [13] can be applied successfully to non-identical bidders only if sufficient qualitative information about bidder asymmetry is publicly known. They proposed a model of *ordered bidders*. Earlier bidders are in some sense expected to have higher valuations. This information could be derived from, for example, zip codes, eBay bidding histories, credit history, previous transactions with the seller, and so on. Leonardi and Roughgarden [17] defined the *monotone price benchmark* $\mathcal{M}^{(2)}(\mathbf{b})$ for every bid vector \mathbf{b} as the maximum revenue obtainable via a monotone price vector — meaning prices are nonincreasing in the bidder ordering — in which every price is at most the second-highest bid.² The value of this benchmark is always at least that of the fixed-price benchmark $\mathcal{F}^{(2)}$, and can be a factor of $\Theta(\log n)$ larger, where n is the number of bidders. Essentially by construction, a digital goods auction that always has revenue at least a constant fraction of $\mathcal{M}^{(2)}$ is simultaneously near-optimal in every Bayesian environment with ordered distributions (where monopoly prices are nonincreasing in the bidder ordering), or when the valuation distribution of each bidder stochastically dominates that of the next one in the ordering (see [17] for details). Examples include uniform distributions with intervals $[0, h_i]$ and nonincreasing h_i 's; exponential distributions with nondecreasing rates; Gaussian distributions with nonincreasing means; and so on. The main result in [17] is a prior-free digital goods auction with ordered bidders with expected revenue $\Omega(\mathcal{M}^{(2)}(\mathbf{b})/\log^* n)$ for every input \mathbf{b} , where n is the number of bidders and $\log^* n$ denotes the number of times that the \log_2 operator can be applied to n before the result drops below a fixed constant.³

1.1 Our Results

We give the first digital goods auction that is $O(1)$ -competitive with the monotone price benchmark $\mathcal{M}^{(2)}$. Our auction is simple and natural. It follows the standard approach of randomly partitioning the bidders into two groups, using one group of bidders to set prices for the other. We restrict prices to be (essentially) all powers of a certain constant, but otherwise our prices are simply the optimal monotone ones for the first bidder group. Finally, to handle inputs where the monotone price benchmark derives most

²This benchmark was also considered earlier, with a different motivation and application, by Aggarwal and Hartline [1].

³Aggarwal and Hartline [1] previously obtained an incomparable guarantee of $\Omega(\mathcal{M}^{(2)}(\mathbf{b})) - O(h \log \log \log h)$, where h is the ratio between the maximum and minimum bids.

of its revenue from a small number of bidders, with constant probability we invoke an auction that is $O(1)$ -competitive with the fixed-price benchmark $\mathcal{F}^{(2)}$.

We extend our results to multi-unit auctions, where the number of items k can be less than the number of bidders. We consider the analog $\mathcal{M}^{(2,k)}$ of the monotone price benchmark, which maximizes only over (monotone) price vectors that sell at most k units. We prove that every auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}$ implies simultaneously near-optimal for a range of Bayesian multi-unit environments — roughly, those in which the (ironed) virtual valuation functions of the bidders form a pointwise total ordering. We also give a general reduction, showing how to build a limited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2,k)}$ from an unlimited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2)}$.

2 Preliminaries

In a *multi-unit auction*, there is one seller, n bidders, and k identical items. Each bidder wants only one good, and has a private — i.e., unknown to the seller — *valuation* v_i . We call the special case where $k = n$ *unlimited supply* or *digital goods*. We study direct-revelation auctions, in which the bidders report bids \mathbf{b} to the seller, and the seller then decides who wins a good and at what price.⁴ For a fixed (randomized) auction, we use $X_i(\mathbf{b})$ and $P_i(\mathbf{b})$ to denote the winning probability and expected payment of bidder i when the bid profile is \mathbf{b} . As in previous works on prior-free auction design, we consider only auctions that are individually rational — meaning $P_i(\mathbf{b}) \leq v_i \cdot X_i(\mathbf{b})$ for every i and \mathbf{b} — and truthful, meaning that for each bidder i and fixed bids \mathbf{b}_{-i} by the other bidders, bidder i maximizes its quasi-linear utility $v_i \cdot X_i(b_i, \mathbf{b}_{-i}) - P_i(b_i, \mathbf{b}_{-i})$ by setting $b_i = v_i$. Since we consider only truthful auctions, from now on we use bids \mathbf{b} and valuations \mathbf{v} interchangeably.

Truthful and individually rational digital goods auctions have a nice canonical form: for every bidder i there is a (possibly randomized) function $t_i(\mathbf{v}_{-i})$ that, given the valuations \mathbf{v}_{-i} of the other bidders, gives bidder i a “take-it-or-leave-it offer” at the price $t_i(\mathbf{v}_{-i})$. This means that bidder i is given a good if and only if $v_i \geq t_i(\mathbf{v}_{-i})$, in which case it is charged the price $t_i(\mathbf{v}_{-i})$. It is clear that every choice (t_1, \dots, t_n) of such functions defines a truthful, individually rational digital goods auction; conversely, every such auction is equivalent to a choice of (t_1, \dots, t_n) [9]. A special case of

⁴For the questions we ask, the “Revelation Principle” (see, e.g., Nisan [20]) ensures that there is no loss of generality by considering only direct-revelation auctions.

such an auction is a *price vector* p , in which each t_i is the constant function $t_i(\mathbf{v}_{-i}) = p_i$. When the supply is limited (i.e., there are $k < n$ copies of the good), truthful auctions induce functions (t_1, \dots, t_n) with the property that, on every input, at most k bidders win.

The *revenue* of an auction on the valuation profile \mathbf{v} is the sum of the payments collected from the winners. Let $v^{(2)}$ denote the second-highest valuation of a profile \mathbf{v} . The *fixed-price benchmark* $\mathcal{F}^{(2)}$ is defined, for each valuation profile \mathbf{v} , as the maximum revenue that can be obtained from a constant price vector whose price is at most $v^{(2)}$:

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{p \leq v^{(2)}} \left(\sum_{i: v_i \geq p} p \right).$$

Now suppose there is a known ordering on the bidders, say $1, 2, \dots, n$. The *monotone-price benchmark* $\mathcal{M}^{(2)}$ is defined analogously to $\mathcal{F}^{(2)}$, except that non-constant monotone price vectors are also permitted:

$$\mathcal{M}^{(2)}(\mathbf{v}) = \max_{v^{(2)} \geq p_1 \geq p_2 \geq \dots \geq p_n} \left(\sum_{i: v_i \geq p_i} p_i \right). \quad (1)$$

Clearly, $\mathcal{M}^{(2)}(\mathbf{v}) \geq \mathcal{F}^{(2)}(\mathbf{v})$ for every input \mathbf{v} .

The monotonicity and upper-bound constraints are enforced only in the computation of the benchmark $\mathcal{M}^{(2)}$. Auctions, while obviously not privy to the private valuations, can employ whatever prices they see fit. This is natural for prior-free auctions and also necessary for non-trivial results [8].

Finally, when we say that an auction is α -*competitive* with or has *approximation factor* α for a benchmark, we mean that the auction's expected revenue is at least a $1/\alpha$ fraction of the benchmark for every input \mathbf{v} .

3 Optimal price scaling

In this section, we propose and analyze a simple auction, which we call OPTIMAL PRICE SCALING (OPS). The basic idea is very natural: partition the sequence of bidders into two random sets, find the optimal $\mathcal{M}^{(2)}$ price vector p for one part A (the *training part*) and offer p to the other part B (the *test part*). We add two twists to this standard algorithmic scheme. First, we restrict the prices to be on discrete levels, so that at every level the prices remain the same, while the price drops from level to level by a constant factor w , where w is a parameter of the algorithm (a particular

value that works is $w = 25$)⁵. The levels of the prices start at the second maximum value $\mathbf{v}^{A(2)}$ of the training part A . The second twist is that we run the algorithm with the dropping prices only with some constant probability and, with the remaining probability, we run a standard digital goods auction with constant competitive ratio. The precise description of the auction is given in Figure 1.

Input: A valuation profile \mathbf{v} for a totally ordered set $N = \{1, 2, \dots, n\}$ of bidders.

1. With probability $1/2$, run a digital goods auction on \mathbf{v} that is $O(1)$ -competitive with respect to the benchmark $\mathcal{F}^{(2)}$. With the remaining probability, run the following steps.
2. Choose a subset $A \subseteq N$ uniformly at random, and partition N into the two sets A and B . Let \mathbf{v}^A denote the valuation profile \mathbf{v} in which we set the values not in A to 0, that is,

$$\mathbf{v}_j^A = \begin{cases} \mathbf{v}_j & \text{if } j \in A \\ 0 & \text{otherwise.} \end{cases}$$

Define \mathbf{v}^B in a similar way. All three sequences \mathbf{v} , \mathbf{v}^A , and \mathbf{v}^B have the same length.

3. Compute an optimal monotone $\mathcal{M}^{(2)}$ price vector p for A with prices restricted to discrete values in $\{\mathbf{v}^{A(2)}/w^k : k = 0, \dots\}$.
4. Sell items to bidders in B only, applying prices p to \mathbf{v}^B .

Figure 1: The auction OPTIMAL PRICE SCALING (OPS).

We next elaborate on the steps of the auction. In the first step, we run an arbitrary digital goods auction that is $O(1)$ -competitive with respect to the fixed-price benchmark $\mathcal{F}^{(2)}$. The best-known approximation factor is 3.12 [16]; there are also very simple auctions with approximation factors 4 [9] and 4.68 [2]. Intuitively, this step is meant to extract good revenue from the set of bidders with valuations almost as high as the second-highest

⁵From our analysis of this algorithm, it follows if we, instead of fixing the prices at discrete levels, offered to the trial part the optimal price of the training part reduced by a factor w , we will still get a truthful algorithm with constant approximation ratio.

valuation.

The second step of the algorithm randomly partitions the bidders into a “training set” A and a “test set” B . Almost all prior-free auctions have this structure, with the bidders in the training set setting prices for those in the test set. We want to keep the three sequences \mathbf{v}^A , and \mathbf{v}^B aligned to simplify the pricing of the next two steps. To do this, we keep all the elements of \mathbf{v} in \mathbf{v}^A and \mathbf{v}^B , but in \mathbf{v}^A we lower the elements that are not in A to 0, and similarly, we lower the elements of sequence \mathbf{v}^B that are not in B to 0.

Also for simplicity, we sell (in the fourth step) only to bidders in the test set B . An obvious optimization is to sell simultaneously to bidders in A , using the bids of B ; this would improve the hidden constant in our approximation guarantee by a factor of 2.

The optimal monotone price vector is the one that maximizes the $\mathcal{M}^{(2)}$ revenue obtained from the bidders in A when prices are scaled down to the next level $\frac{\mathbf{v}^{A(2)}}{w^k}$. The scaling down of the prices won’t reduce the optimal revenue on the A set of bidders for more than a factor of w . The final step applies the prices to bidders in the test set B .

The OPS auction is truthful, as each bidder faces a take-it-or-leave-it offer at a price that is independent of its reported valuation. We also note that the OPS auction can be implemented in polynomial time, as p can be computed efficiently using dynamic programming.

Our main result is a prior-free approximation guarantee for the OPS auction.

Theorem 3.1 *There is a constant $c_0 > 0$ such that, for every valuation profile \mathbf{v} , the expected revenue of the OPS auction is at least $c_0 \cdot \mathcal{M}^{(2)}(\mathbf{v})$.*

We outline the main ideas of the proof of the theorem here and we present the details below. Let us first make a useful assumption

Definition 3.2 (First running assumption) *In the analysis of the algorithm, we will assume that the two bidders of \mathbf{v} with the highest valuation fall into A . Therefore*

$$\mathbf{v}^{(2)} = \mathbf{v}^{A(2)}.$$

This assumption holds with probability $1/4$, which will essentially increase the competitive ratio of the analysis below by a factor of 4. We remove this assumption, when we put all the pieces of the analysis together.

Let us now fix some notation:

Definition 3.3 Define the k -th price level $p_{(k)}$ to be

$$p_{(k)} = \mathbf{v}^{(2)} / w^k.$$

Let also J_k denote the interval in which the prices of the auction are at the k -th level:

$$J_k = \{j : p_j = p_{(k)}\}.$$

The intervals J_0, \dots, J_m are defined in the third step of the algorithm from \mathbf{v}^A . If there are many optimal solutions in the third step, we fix one and we use it in the final step of the algorithm.

The intuition for the algorithm is that the optimal revenue for the whole sequence \mathbf{v} is up to a constant factor equal to the expected optimal revenue from the subsequence \mathbf{v}^A . We then need only to compare the optimal expected revenues from the subsequences \mathbf{v}^A and \mathbf{v}^B . If there were many values in a pricing level J_k , then the random partition is expected to split almost evenly the high-valued bids between the training set A and the test set B and this will allow us to relate their revenue. One problem with this approach is that some levels may have few items; another more subtle problem is that even when a level has many items, we cannot easily argue that the two parts have almost the same number of high values, because there is a bias towards A by the way the levels were created.

To resolve both these issues and in order to compare the revenue of A and B , we use as an intermediary a set of some fixed intervals, defined with respect to the set \mathbf{v} of all values (in contrast to the way that the levels are created by the pricing of A). The set of fixed intervals, which we will call *primal intervals*, consists of one collection of intervals for each pricing level. We want them to have two crucial properties: first, *every primal interval* to have almost the same fraction (and in particular a fraction in $[1/3, 2/3]$) in A and B of the bids higher than the pricing level; and second, that the revenue of A is captured up to a constant factor by a set of primal intervals.

3.1 Primal intervals

We now define the set of primal intervals. They will be defined in terms of \mathbf{v} , the whole sequence of values. We will need some definitions first:

Definition 3.4

- Define N_ℓ to be the set of values greater or equal to the ℓ -th price level.

$$N_\ell = \{j : \mathbf{v}_j \geq p_{(\ell)}\} \tag{2}$$

- Define $\mu_\ell(I)$, the mass of an interval I at price level ℓ , to be the number of its values that are in N_ℓ (at price level ℓ or higher).

$$\mu_\ell(I) = |I \cap N_\ell|.$$

For each price level ℓ , we create a set of intervals \mathcal{I}_ℓ which we will call *primal intervals at price level ℓ* . When we refer to the mass of a primal interval $I \in \mathcal{I}_\ell$, we will mean $\mu_\ell(I)$, its mass at its price level. We will denote by $\text{start}(I)$ and $\text{end}(I)$ the beginning and end of an interval.

The primal intervals are defined with respect to 3 parameters: the parameter w which defines the price levels in the algorithm, a parameter $\epsilon \in (0, 1)$, which will be fixed later (a value of $\epsilon = 1/25$ will work for the proof) and a parameter s_0 ; notice that the parameters ϵ and s_0 are used only in the analysis and they are not part of the algorithm.

We begin the construction with a fixed interval $I_o \in \mathcal{I}_0$ which starts at the first value and has $\mu_0(I_o) = s_0$, and we define the other intervals recursively. If $I \in \mathcal{I}_\ell$ is an interval that has already been defined, then we add the following 5 intervals every one of which has mass $(1 + \epsilon) \cdot \mu_\ell(I)$; if there is not enough mass to create some of these 5 intervals, we still create them with the maximum possible mass to facilitate the recursive construction, we call them incomplete, and we throw them away at the end of the construction.

1. An interval at level \mathcal{I}_ℓ which starts at $\text{start}(I)$.
2. An interval at level \mathcal{I}_ℓ which starts at $\text{end}(I)$.
3. An interval at level $\mathcal{I}_{\ell+1}$ which starts at $\text{start}(I)$.
4. An interval at level $\mathcal{I}_{\ell+1}$ which starts at $\text{end}(I)$.
5. An interval at level \mathcal{I}_ℓ which *ends* at $\text{end}(I)$.

Figure 2 illustrates the construction. The reason for keeping around incomplete intervals is to allow an interval at level ℓ to be able to create intervals at all higher levels. Specifically, we want an interval $I \in \mathcal{I}_\ell$ to create every possible interval at level $\ell + r$, $r \geq 1$, that start at the end (or beginning) of I and have mass $(1 + \epsilon)^r$ times the mass of I . For instance, even if the interval at level $\ell + r$ is incomplete, we keep it in the recursive construction to produce the potentially complete intervals at level $\ell + r + 1$ by using Rules 3 and 4. We extend the idea of keeping around incomplete intervals to the original interval I_o : if there are not enough elements with high mass,




Figure 2: The 5 intervals created from interval I . Rules 1, 2 and 5 create intervals at the same level with I , while Rules 3 and 4 create intervals at the next level. The drawing indicates the mass of the intervals—equal to $(1 + \epsilon)$ the mass of I —not their actual length.

we define I_o to include the whole sequence, we call it incomplete, we use it to create recursively the other intervals, and we throw it away at the end. The mass of an incomplete interval in the construction is assumed to be the originally intended mass. For example, suppose that a complete interval $I \in \mathcal{I}_\ell$ creates an incomplete interval I' which in turn creates a complete interval I'' ; the mass of I'' is $(1 + \epsilon)^2$ times the mass of I , irrespectively of the actual mass of I' .

We now show a crucial property of the primal intervals: with probability almost 1, every primal interval contains almost equal number of elements in A and B . More precisely, let us call a primal interval $I \in \mathcal{I}_\ell$ *balanced*, if at least a third of the elements in $I \cap N_\ell$ belong to A and at least another third belongs to B . For every ϵ , by selecting the parameter s_0 appropriately, we can make the probability of having some unbalanced primal interval arbitrarily small. Before we state and prove a useful fact:

Lemma 3.5 *If d is a positive constant such that $d^{\epsilon \cdot s_0} \leq 1/10$, then*

$$\sum_{\ell} \sum_{I \in \mathcal{I}_\ell} d^{\mu_\ell(I)} \leq 2 \cdot d^{s_0}.$$

Proof: Suppose that the initial interval in the construction has mass $s \geq s_0$ and define

$$f(s) = \sum_{\ell} \sum_{I \in \mathcal{I}_\ell} d^{\mu_\ell(I)},$$

to be the maximum value attained from the intervals produced from s by any sequence of values. Essentially, we want to bound $f(s_0)$. By the recursive construction of the intervals, we get that

$$f(s) \leq c^s + 5f((1 + \epsilon) \cdot s).$$

The lemma now follows from showing $f(s) \leq 2d^s$ by backwards induction. For $s > n$, this is vacuously true since there are no intervals. Suppose then that $f((1 + \epsilon)s) \leq 2d^{(1+\epsilon)s}$, and we bound

$$\begin{aligned} f(s) &\leq d^s + 5f((1 + \epsilon)s) \leq d^s + 10d^{(1+\epsilon)s} \\ &= d^s + 10d^s d^{\epsilon \cdot s} \leq d^s + 10d^s d^{\epsilon s_0} \leq d^s + 10\frac{1}{10}d^s = 2d^s. \end{aligned}$$

■

Lemma 3.6 *There is a large enough constant s_0 , such that for every \mathbf{v} and ϵ , with probability at least $31/32$ every primal interval is balanced.*

Proof: The probability that a primal interval I is *not* balanced is at most $d^{\mu_\ell(I)}$ for some constant d (Chernoff bound). Therefore, the probability that some interval is not balanced is at most

$$\sum_{\ell} \sum_{I \in \mathcal{I}_\ell} d^{\mu_\ell(I)} \leq 2d^{s_0},$$

which drops exponentially in s_0 . By selecting s_0 large enough, we can make this probability as small as desired. ■

3.2 Matching intervals

To gain some intuition of the use of the primal intervals, let us assume that we could associate with every interval J_k a primal interval $I_k \in \mathcal{I}_k$ which is also a subset of J_k . Let us further assume that among the possible I_k 's, we always select the interval with maximum length. There are two facts that allow us to relate the revenues from B and from A : the first is that by the construction of the primal intervals, the length of I_k contains at least $1/(1 + 2\epsilon)$ of J_k ; the second is that the interval I_k is balanced. The latter shows that the revenue from B in I_k is a constant fraction of the revenue from N in I_k , which in turn is a constant fraction from the revenue from A in J_k . The problem is that there may be intervals J_k that contain no primal interval in \mathcal{I}_k . To deal with this problem, we use a charging argument which shows that the revenue lost from such intervals is not significant.

To turn this idea into a concrete proof, we give a process that attempts to match every interval J_k to some primal interval in \mathcal{I}_k ; if the process succeeds, we will call J_k *matched* and we will denote its matching primal interval by I_k ; if the process does not succeed for some interval J_k we will call J_k *unmatched*. We first match every interval J_k that contains some primal interval at the same price level.

Definition 3.7 *An interval J_k is called nice, if there is an primal interval in $I \in \mathcal{I}_k$ such that $\text{start}(J_k) \leq \text{start}(I) \leq \text{end}(I) \leq \text{end}(J_k)$. We match J_k with the maximum such interval I and we call it I_k .*

We process the remaining intervals from left to right. To bootstrap the process, we make the assumption that the first interval J_0 is nice; we will revisit and remove this assumption at the end of the proof.

Definition 3.8 (Second running assumption) *Interval J_0 is nice, i.e., it contains a primal interval of \mathcal{I}_0 .*

Starting with each nice interval, we process its following intervals from left to right. We attempt to find a matching interval by creating a sequence of candidate intervals. The definition of the sequence of candidate matching intervals will be also useful later in the presentation of the main argument. Figure 3 illustrates the process which is defined in detail below:

Definition 3.9 (Matching) *For every interval J_t , we associate a non-empty sequence of candidate matching intervals, all of which are at price level t . If J_t is a nice interval, the sequence contains only the matching interval I_t . For every other interval J_t , the sequence of candidate matching intervals is created as follows:*

The first candidate matching interval will be called I'_t plays also a central role in the proof and it is based on the candidate matching intervals of the previous interval J_{t-1} . In particular, if the previous interval is matched, then I'_t is the interval produced from I_{t-1} by Rule 4. Otherwise, if the previous interval is unmatched, I'_t is the interval produced from I'_{t-1} by Rule 3.

If $\text{end}(I'_t) > \text{end}(J_t)$, the sequence of candidate matching intervals contains only I'_t and we leave J_t unmatched.

Otherwise, the sequence of candidate matching intervals is produced from this first interval I'_t by repeated applications of Rule 1, as long as the produced intervals do not extend beyond J_t (i.e., they end at or before $\text{end}(J_t)$).

We match J_t with the maximum interval in the candidate matching sequence, which we name I_t .

3.3 Main argument

The set of nice intervals play central role in relating the revenue of A and B . Before we proceed with the central argument, we give some essential definitions and prove some useful facts about the matched intervals.




Figure 3: An example of the matching process. Interval J_r is nice. Interval I'_{r+1} is created from I_r by Rule 4. It extends beyond J_{r+1} , so we leave J_{r+1} unmatched. Interval I'_{r+2} is created from I'_{r+1} by Rule 3. Similarly, we leave J_{r+2} unmatched. Interval I'_{r+3} is created from I'_{r+2} by Rule 3; it produces the sequence of candidate intervals $I'_{r+3}, I''_{r+3}, \dots, I_{r+3}$ by repeated applications of Rule 1. Interval J_{r+3} is matched to I_{r+3} . Intervals J_{r+4} and I_{r+4} exhibit a case in which the matched intervals do not have to intersect.

Definition 3.10 Let $R(p, v)$ denote the total revenue extracted from the sequence of values v with prices p . Let $R_k(J, v)$ denote the revenue that can be extracted from an interval J of v to which we offer the k -level price $p_{(k)}$: $R_k(J, v) = p_{(k)} \cdot |\{i : i \in J \wedge v_i \geq p_{(k)}\}|$. In particular, $R_k(J, \mathbf{v})$ is equal to the mass of J multiplied by $p_{(k)}$.

Lemma 3.11 The matching algorithm has the following properties:

1. If J_t is nice, then $\mu_t(J_t) \leq (1 + 2\epsilon)\mu_t(I_t)$.
2. For every constant $c_3 > 2$ let

$$c_4 = \frac{c_3}{c_3 - 2} \cdot (2 + \epsilon) \cdot \epsilon.$$

If J_t is not nice, then at least one of the following holds

$$\mu_t(J_t) \leq c_3 \cdot \mu_t(I'_t) \tag{3}$$

$$\mu_t(J_t) \leq c_4 \cdot \mu_t(I_t) \tag{4}$$

Proof: The properties follow directly from the definitions and the process. For example, the first property, follows by the maximality of I_t taking into account Rules 1 and 5; this is the only place in the proof where we use Rule 5.




Figure 4: Illustrating the proof of Lemma 3.11.

The second property is more involved and uses the fact that J_t is not nice. Notice first that if J_t is not matched, then $\mu_t(J_t) \leq \mu_t(I'_t) \leq c_3 \mu_t(I'_t)$ and the property holds. So now assume that J_t is matched and consider the sequence of intervals K_0, K_1, \dots, K_{r+1} , produced when we start with $K_0 = I'_t$ and repeatedly apply Rule 2; the sequence ends when the produced interval extends beyond J_t ; that is, $\text{start}(K_{r+1}) < \text{end}(J_t) < \text{end}(K_{r+1})$ (See Figure 4).

Since J_t is not nice, it does not include any of these intervals and in particular interval K_r , so we must have that $\mu_t(J_t) \leq \mu_t(K_r) + \mu_t(K_{r+1})$. This bounds the mass of J_t by

$$\mu_t(J_t) \leq ((1 + \epsilon)^r + (1 + \epsilon)^{r+1}) \mu_t(I'_t) = (2 + \epsilon)(1 + \epsilon)^r \mu_t(I'_t).$$

On the other hand, I_t must contain almost all the mass of intervals K_0, \dots, K_r . In particular, by the maximality of I_t , if we apply Rule 1 to I_t , we get an interval which extends beyond J_t , and it certainly contains the intervals K_0, \dots, K_r . Thus the mass of I_t is at least

$$\begin{aligned} \mu_t(I_t) &\geq (\mu_t(K_0) + \dots + \mu_t(K_r)) / (1 + \epsilon) \\ &= (1 + (1 + \epsilon) + \dots + (1 + \epsilon)^r) / (1 + \epsilon) \cdot \mu_t(I'_t) \\ &= ((1 + \epsilon)^{r+1} - 1) / (\epsilon(1 + \epsilon)) \cdot \mu_t(I'_t). \end{aligned}$$

The lemma now follows from the above two bounds. Specifically, if $\mu_t(J_t) > c_3 \mu_t(I'_t)$, then $c_3 \geq (2 + \epsilon)(1 + \epsilon)^r$. We can then bound

$$\begin{aligned} \mu_t(J_t) &\leq (2 + \epsilon)(1 + \epsilon)^r \mu_t(I'_t) \\ &\leq (2 + \epsilon)(1 + \epsilon)^r \frac{\epsilon(1 + \epsilon)}{(1 + \epsilon)^{r+1} - 1} \mu_t(I_t) \\ &\leq \frac{c_3}{c_3 - 2} \cdot (2 + \epsilon) \cdot \epsilon \cdot \mu_t(I_t). \end{aligned}$$

■

The essential part of proving the constant approximation ratio is a charging argument. The idea is to charge each of the values $R_t(J_t, \mathbf{v})$ to some

value $R_s(I_s, \mathbf{v})$ of some nice interval J_s . To do this, we process the intervals from left to right in phases. A phase begins with a nice interval and ends exactly before the next nice interval (or at the last interval).

Lemma 3.12 (Phase) *Let $s, \dots, s' - 1$ be a phase, that is, J_s is a nice interval and none of the intervals J_r for $r = s + 1, \dots, s' - 1$ is nice. Then there is a constant c such that*

$$\sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v}) \leq c \cdot R_s(I_s, \mathbf{v}) \quad (5)$$

Proof: We charge each interval J_t , $t = s + 1, \dots, s' - 1$ to some primal interval, according to the cases of the previous lemma, as follows:

1. If $\mu_t(J_t) \leq c_3 \cdot \mu_t(I'_t)$, we charge it to the previous matched interval. Specifically, if $\sigma(t)$ denotes the previous matched interval (for example, in Figure 3, $\sigma(r + 2) = r$), then

$$R_t(J_t, \mathbf{v}) \leq c_3 R_t(I'_t, \mathbf{v}) \leq c_3 \left(\frac{1 + \epsilon}{w} \right)^{t - \sigma(t)} R_{\sigma(t)}(I_{\sigma(t)}, \mathbf{v}).$$

2. Otherwise we charge J_t to I_t . Specifically, by Lemma 3.11 we have

$$R_t(J_t, \mathbf{v}) \leq c_4 R_t(I_t, \mathbf{v}).$$

If we sum the above for $t = s + 1, \dots, s' - 1$, we get

$$\sum_{t=s+1}^{s'-1} R_t(J_t, \mathbf{v}) \leq \sum_{t=s}^{s'-1} \alpha_t \cdot R_t(I_t, \mathbf{v}), \quad (6)$$

for some α_t which is the sum of all the above charges. The crucial property of the parameters is that the total charge to every interval is low (less than $1/3$) in every possible scenario. Specifically, the total charge to α_t can be at most c_4 (Case 2), plus at most $c_3 \cdot \sum_{i=1}^{\infty} \left(\frac{1+\epsilon}{w} \right)^i = c_3 \cdot \frac{1+\epsilon}{w-1-\epsilon}$ from Case 1. The total charge is

$$\alpha_t \leq \alpha = c_4 + c_3 \cdot \frac{1 + \epsilon}{w - 1 - \epsilon}, \quad (7)$$

which is less than $1/3$ if we select the parameters appropriately.

We now use the optimality of price vector p to bound $\sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v})$. Consider decreasing the prices in the interval $\cup_{t=s}^{s'-1} J_t$ to level $s' - 1$. By the

optimality of the price profile p , this cannot give better revenue for \mathbf{v}^A , and therefore

$$\sum_{t=s}^{s'-1} R_t(I_t, \mathbf{v}^A) \leq \sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v}^A). \quad (8)$$

We want now to combine (6) and (8), but one involves \mathbf{v} and the other \mathbf{v}^A , so we employ the balance property to transform (8) to:

$$\sum_{t=s}^{s'-1} R_t(I_t, \mathbf{v}) \leq 3 \sum_{t=s}^{s'-1} R_t(I_t, \mathbf{v}^A) \leq 3 \sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v}^A) \leq 3 \sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v}). \quad (9)$$

We will also need the first property of Lemma 3.11:

$$R_s(J_s, \mathbf{v}) \leq (1 + 2\epsilon) R_s(I_s, \mathbf{v}). \quad (10)$$

If we now combine Equations (6), (9), and (10) (multiply (9) by α , take their sum, and divide the result by $1 - 3\alpha$), we get

$$\sum_{t=s}^{s'-1} R_t(J_t, \mathbf{v}) \leq \frac{1 + 2\epsilon}{1 - 3\alpha} R_s(J_s, \mathbf{v})$$

(this is where it is crucial that α is less than $1/3$). Letting now $c = \frac{1+2\epsilon}{1-3\alpha}$, we get the lemma. ■

By simply summing the inequalities of the previous lemma for all nice intervals, we get:

Lemma 3.13 *Let $J_{q_0}, \dots, J_{q_{m'}}$ be the nice intervals. Assuming that the first interval is matched (Definition 3.8), or equivalently that $q_0 = 0$, there is a constant c such that*

$$\sum_{j=0}^m R_j(J_j, \mathbf{v}) \leq c \cdot \sum_{k=0}^{m'} R_{q_k}(I_{q_k}, \mathbf{v}). \quad (11)$$

The previous lemma relates the revenue from the whole sequence to the revenue from the matching primal intervals of nice intervals. Because on every primal balanced interval the revenues of \mathbf{v} , \mathbf{v}^A , and \mathbf{v}^B are all within a constant factor, we can immediately relate $R(p, \mathbf{v}^A)$ to $R(p, \mathbf{v}^B)$.

Lemma 3.14 *With the assumptions of the previous lemma,*

$$R(p, \mathbf{v}^A) \leq 3c \cdot R(p, \mathbf{v}^B).$$




Figure 5: An example of possible charging. J_r is a nice interval. J_{r+1} is charged either to I_r (if the mass of I'_{r+1} is small) or to I_{r+1} (otherwise). Similarly J_{r+2} is charged either to I_{r+1} or to I_{r+2} depending on the mass of I'_{r+2} (shown). J_{r+3} is charged to I_{r+2} (because it is not matched). J_{r+4} is charged either to I_{r+2} or to I_{r+4} , depending on the mass of I'_{r+4} .

Proof: We have the following derivations

$$\begin{aligned}
R(p, \mathbf{v}^A) &= \sum_{j=0}^m R_j(J_j, \mathbf{v}^A) \\
&\leq \sum_{j=0}^m R_j(J_j, \mathbf{v}) \\
&\leq c \cdot \sum_{k=0}^{m'} R_{q_k}(I_{q_k}, \mathbf{v}) \\
&\leq 3c \cdot \sum_{k=0}^{m'} R_{q_k}(I_{q_k}, \mathbf{v}^B) \\
&\leq 3c \cdot \sum_{k=0}^{m'} R_{q_k}(J_{q_k}, \mathbf{v}^B) \\
&\leq 3c \cdot \sum_{j=0}^m R_j(J_j, \mathbf{v}^B) \\
&\leq 3c \cdot R(p, \mathbf{v}^B)
\end{aligned}$$

The first inequality is trivial because \mathbf{v}^A has smaller values than \mathbf{v} , the second inequality comes from the main lemma (Lemma 3.13), the third inequality comes from the balance condition, the fourth inequality is based

on the fact that I_{q_k} is contained in J_{q_k} . ■

3.4 Stitching everything together

In this subsection, we revisit and remove the First Running Assumption. For simplicity, we keep the Second Running Assumption and we deal with it later.

Lemma 3.15 *For every valuation profile \mathbf{v} ,*

$$\Pr \left[R(p, \mathbf{v}^A) \geq \frac{1}{3w} \cdot \mathcal{M}^{(2)}(\mathbf{v}) \right] \geq \frac{1}{16}.$$

Proof: Let p^* be the optimal prices for $\mathcal{M}^{(2)}(\mathbf{v})$. With probability $1/4$, the bidders with the highest and second-highest valuations lie in A , i.e., $\mathbf{v}^{(2)} = \mathbf{v}^{A(2)}$. Given this event, the conditional expected revenue from bidders in A and B under the price vector p^* is at least $\mathcal{M}^{(2)}(\mathbf{v})/2$ and at most $\mathcal{M}^{(2)}(\mathbf{v})/2$, respectively. By Markov's inequality, the conditional expected revenue from bidders in A under p^* is at least $\frac{1}{3}\mathcal{M}^{(2)}(\mathbf{v})$ with probability at least $\frac{1}{4}$.

The optimal revenue $R(p, \mathbf{v}^A)$ from A cannot be worse than the revenue from the pricing vector p^* scaled down to the next power $\frac{\mathbf{v}^{(2)}}{w^k}$. Since with this scaling down, we loose at most a factor of w in the revenue, we conclude that with probability at least $\frac{1}{16}$, $R(p, \mathbf{v}^A) \geq \mathcal{M}^{(2)}(\mathbf{v})/3w$. ■

If we choose appropriately the constant s_0 , Lemma 3.6 ensures that for every constant $\epsilon \in (0, 1]$, and every valuation profile \mathbf{v} , every primal interval is balanced with probability at least $31/32$. (Note that for small n , this is vacuously true since there are no primal intervals).

The probability that all positive events (that the two higher values are in A , that $\mathcal{M}^{(2)}(\mathbf{v}^A) \geq \frac{1}{3w} \cdot \mathcal{M}^{(2)}(\mathbf{v})$, and that all primal intervals are balanced) hold is at least $1 - (\frac{1}{32} + \frac{15}{16}) = \frac{1}{32}$. So, in the analysis of the performance of the OPTIMAL PRICE SCALING (OPS) auction, we can assume that these events hold, which will increase the approximation ratio by a factor of 32.

By Lemma 3.14, we have that there are parameters ϵ , w and s_0 such that

$$R(p, \mathbf{v}^B) \geq \frac{1}{3c} R(p, \mathbf{v}^A) \geq \frac{1}{3c} \frac{1}{3w} \mathcal{M}^{(2)}(\mathbf{v}) = c_1 \cdot \mathcal{M}^{(2)}(\mathbf{v}),$$

(where $c_1 = 1/(9cw)$), which establishes the theorem under the assumption that the first interval J_0 is nice.

3.5 Padding

In the previous subsection, we analyzed the algorithm and bounded its revenue under the assumption that the first interval J_0 is nice, i.e., it contains a primal interval. We now show how to remove this assumption by padding \mathbf{v} on the left with a sequence of $\mathbf{v}^{(2)}$ values which will guarantee that the assumption holds. Then it suffices to show that the padding does not affect the revenue significantly.

In particular, let $\hat{\mathbf{v}}$ be the sequence consisting of $s'_0 = 10s_0$ values $\mathbf{v}^{(2)}$ followed by \mathbf{v} . With probability very close to 1, at least s_0 of these values go to $\hat{\mathbf{v}}^A$, and therefore the interval J_0 contains the initial primal interval⁶. Notice that $\mathcal{M}^{(2)}(\hat{\mathbf{v}}) = \mathcal{M}^{(2)}(\mathbf{v}) + s'_0 \cdot \mathbf{v}^{(2)}$ because every optimal decreasing pricing can sell to the first s_0 bidders at price $\mathbf{v}^{(2)}$ without affecting the pricing for the remaining bidders; this is because the prices for the remaining bidders cannot exceed $\mathbf{v}^{(2)}$. Similar constraints hold for $\mathcal{M}^{(2)}(\hat{\mathbf{v}}^A)$ and $\mathcal{M}^{(2)}(\hat{\mathbf{v}}^B)$.

Let ALG_M be the expected revenue when we apply the pricing scheme to \mathbf{v}^A and \mathbf{v}^B and let ALG_F be the expected revenue when we apply an algorithm which is competitive against $\mathcal{F}^{(2)}$. The expected revenue of the algorithm is $(\text{ALG}_M + \text{ALG}_F)/2$. The revenue extracted from the part against $\mathcal{F}^{(2)}$ will play a crucial role in relating the revenue of the original sequence with the revenue with the padded sequence. For the padded sequence, let $\hat{\text{ALG}}_M$ be the corresponding expected revenue for the padded sequence, that is, when we apply the pricing scheme to $\mathbf{v}^{\hat{A}}$ and $\mathbf{v}^{\hat{B}}$. From the previous subsection, we have that

$$\hat{\text{ALG}}_M \geq c_1 \cdot \left(\mathcal{M}^{(2)}(\mathbf{v}) + s'_0 \cdot \mathbf{v}^{(2)} \right),$$

for some constant c_1 . Also, it is easy to see that the prices of the algorithm applied to $\hat{\mathbf{v}}$ match the prices of the algorithm applied to \mathbf{v} , so we have that $\hat{\text{ALG}}_M \leq \text{ALG}_M + s'_0 \cdot \mathbf{v}^{(2)}$. Putting the last two inequalities together, we get

$$\begin{aligned} \text{ALG}_M &\geq c_1 \cdot \left(\mathcal{M}^{(2)}(\mathbf{v}) + s'_0 \cdot \mathbf{v}^{(2)} \right) - s'_0 \cdot \mathbf{v}^{(2)} \\ &= c_1 \cdot \mathcal{M}^{(2)}(\mathbf{v}) - (1 - c_1) \cdot s'_0 \cdot \mathbf{v}^{(2)}. \end{aligned}$$

The second term is the reason for targeting both $\mathcal{F}^{(2)}$ and $\mathcal{M}^{(2)}$ and explains the need to run a competitive algorithm against $\mathcal{F}^{(2)}$ with probability $1/2$.

⁶We need to select s'_0 large enough so that this probability is greater than $1 - 1/32$ to guarantee that all assumptions and events of the previous subsection hold with positive probability.

To account for the second term, we observe that with probability $1/2$, the algorithm gets revenue at least $c'_1 \cdot \mathcal{F}^{(2)}(\mathbf{v}) \geq c'_1 \cdot 2\mathbf{v}^{(2)}$, for some constant c'_1 . That is,

$$\text{ALG}_F \geq c'_1 \cdot 2\mathbf{v}^{(2)}.$$

Combining the two bounds for ALG , we get

$$\begin{aligned} \text{ALG} &\geq \frac{1}{2} \text{ALG}_M \\ \text{ALG} &\geq \frac{1}{2} \text{ALG}_F \\ \text{ALG} &\geq \frac{(1 - c_1) \cdot s'_0 \cdot \left(\frac{1}{2} \text{ALG}_F\right) + 2c'_1 \cdot \left(\frac{1}{2} \text{ALG}_M\right)}{(1 - c_1) \cdot s'_0 + 2c'_1} \\ &= \frac{c_1 \cdot c'_1}{(1 - c_1) \cdot s'_0 + 2c'_1} \cdot \mathcal{M}^{(2)}(\mathbf{v}), \end{aligned}$$

which shows that the algorithm has constant approximation ratio.

4 Multi-Unit Auctions

In this section we extend our results to multi-unit auctions with limited supply. To develop this theory, we extend the monotone price benchmark $\mathcal{M}^{(2)}$ to the case of an arbitrary number $k \geq 2$ of units for sale. We call a price vector p *feasible* for the valuation profile \mathbf{v} and supply limit k if: (i) $p_1 \geq p_2 \geq \dots \geq p_n$; (ii) all prices are at most the second-highest valuation of \mathbf{v} ; and (iii) there are at most k bidders i with $v_i > p_i$. We allow our benchmark to break ties in an optimal way. More precisely, the revenue earned by a feasible price vector is $\sum_{i: v_i > p_i} p_i$ plus, if there are ℓ items remaining, the sum of the prices offered to up to ℓ bidders i with $v_i = p_i$. We define the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ as the maximum revenue obtained by a price vector that is feasible for \mathbf{v} and k .

There are two main issues to address. The first issue is to identify a class of priors \mathcal{F}_i such that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is a meaningful benchmark for prior-free approximation, i.e., it simultaneously approximates all optimal auctions in multi-unit Bayesian settings with priors drawn from the class. The challenge, relative to the unlimited-supply setting introduced in [17], is that limited-supply Bayesian optimal auctions exhibit more complex behavior than unlimited-supply ones. Section 4.1 shows, essentially, that the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is meaningful for any valuation distributions that have pointwise ordered ironed virtual valuations. The second issue is to design auctions competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. We accomplish this

through a general reduction, showing how to build a limited-supply auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2,k)}(\mathbf{v})$ from a digital goods auction that is $O(1)$ -competitive w.r.t. $\mathcal{M}^{(2)}$.

4.1 Justifying the k -Unit Monotone Price Benchmark

The goal of this section is to prove that every prior-free auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ has expected revenue at least a constant fraction of optimal in every Bayesian multi-unit environment with valuation distributions lying in a prescribed class. Making this precise requires some terminology and facts from the theory of Bayesian optimal auction design, as developed by Myerson [19]. See also the exposition by Hartline [11].

Consider a bidder with valuation drawn from a prior distribution \mathcal{F} with positive and continuous density f on some interval. The *virtual value* v at a point v in the support is defined as

$$\phi(v) = v - \frac{1 - \mathcal{F}(v)}{f(v)}.$$

For example, if \mathcal{F} is the uniform distribution on $[0, a]$, then the corresponding virtual valuation function is $\phi(v) = 2v - a$.

For clarity, we first discuss the case of *regular* distributions, meaning distributions with nondecreasing virtual valuation functions. In this case, the Bayesian optimal auction awards items to the (at most k) bidders with the highest positive virtual valuations. The payment of a winning bidder is the minimum bid at which it would continue to win (keeping others' bids the same). That is, if the $(k+1)$ th highest virtual valuation is z , then every winning bidder i pays $\phi_i^{-1}(\max\{0, z\})$. For these prices to be related to the monotone price benchmark, we need to impose conditions on the $\phi_i^{-1}(z)$'s. This contrasts with unlimited-supply settings, where restricting the $\phi_i^{-1}(0)$'s — that is, the monopoly reserve prices — to be nonincreasing in i is enough to justify the monotone-price benchmark [17]. Since the $(k+1)$ th highest virtual valuation could be anything, the natural extension of the condition in [17] is to restrict $\phi_i^{-1}(z)$ to be nonincreasing in i for every non-negative number z .

Accommodating irregular distributions, for which the optimal Bayesian auction is more complicated, presents additional complications. Each virtual valuation function ϕ_i is replaced by the “nearest nondecreasing approximation”, called the *ironed virtual valuation function* $\bar{\phi}_i$. The optimal auction awards the items to the (at most k) bidders with the highest positive ironed

virtual valuations. Since ironed virtual valuation functions typically have non-trivial constant regions, ties can occur, and we assume that ties are broken randomly. That is, if there are k items, a group S of bidders with identical ironed virtual values $z > 0$, $\ell < k$ bidders with ironed virtual value greater than z , and $\ell + |S| > k$, then $k - \ell$ winners from S are chosen uniformly at random.

We call valuation distributions $\mathcal{F}_1, \dots, \mathcal{F}_n$ *pointwise ordered* if $\bar{\phi}_i^{-1}(z)$ is nonincreasing in i for every non-negative z .⁷ The motivating parametric examples discussed earlier — uniform distributions with intervals $[0, h_i]$ and nonincreasing h_i 's, exponential distributions with nondecreasing rates, and Gaussian distributions with nonincreasing means — are pointwise ordered in this sense.

We also require a second condition, which we inherit from the standard i.i.d. unlimited-supply setting. The issue is that, with arbitrary irregular distributions, no prior-free auction can be simultaneously near-optimal in all Bayesian environments, even with i.i.d. bidders and unlimited supply.⁸ Various mild conditions are sufficient to rule out this problem; see [13] for a discussion. Here, for simplicity, we restrict attention to *well-behaved* Bayesian multi-unit environments, meaning that the Bayesian optimal auction derives at most a constant fraction (90%, say) of its revenue from outcomes in which some winner is charged a price higher than the second-highest valuation. (Such a winner is necessarily the bidder with the highest valuation.) Standard distributions always yield well-behaved environments. Even pathological distributions produce well-behaved environments provided the market is sufficiently large (e.g., there are enough bidders drawn i.i.d. from each of the distributions).

Our main result in this section is that approximating the k -unit monotone price benchmark guarantees simultaneous approximation of the optimal auction in all well-behaved Bayesian multi-unit environments with pointwise ordered distributions.

Theorem 4.1 *If the expected revenue of the multi-unit auction \mathcal{A} is at least a constant fraction of $\mathcal{M}^{(2,k)}(\mathbf{v})$ on every input, then, in every well-behaved*

⁷Since $\bar{\phi}_i$ is continuous and nondecreasing, $\bar{\phi}_i^{-1}(z)$ is an interval. If the inverse image has multiple points, we define $\bar{\phi}_i^{-1}(z)$ by the infimum. If the inverse image is empty, we define $\bar{\phi}_i^{-1}(z)$ as the left or right endpoint of the distribution's support, as appropriate.

⁸Informally, consider valuation distributions that take on only two values, one very large (say M) and the other 0. Suppose the probability of having a very large valuation is very small (say $1/n^2$). If the distribution is known, the optimal auction uses a reserve price of M for each bidder. Elementary arguments, as in [13], show that no single auction is near-optimal for all values of M .

multi-unit Bayesian environment with pointwise ordered distributions, the expected revenue of \mathcal{A} is at least a constant fraction of that of the optimal auction for the environment.

Proof: Fix an auction that is β -competitive with $\mathcal{M}^{(2,k)}(\mathbf{v})$ on every input. Fix a well-behaved Bayesian multi-unit environment with pointwise ordered valuation distributions $\mathcal{F}_1, \dots, \mathcal{F}_n$. Let \mathcal{A}^* be the optimal auction for this environment. We claim that, for every input \mathbf{v} in which the revenue collected by \mathcal{A}^* from the bidder with the highest valuation is at most the second-highest valuation, the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least half the expected revenue of \mathcal{A}^* on \mathbf{v} . This implies that the expected revenue of \mathcal{A} is at least $1/2\beta$ times that of \mathcal{A}^* on this input. Since the environment is well behaved, the theorem follows.

To prove the claim, fix an input \mathbf{v} , as above. Recall that \mathcal{A}^* , as a Bayesian optimal auction, awards items to the (at most k) bidders with the highest positive ironed virtual valuations, breaking ties randomly. The tricky case of the proof is when ties occur. Assume there are k items, a group S of bidders with common ironed virtual value $z > 0$, and a group T of $\ell \in (k - |S|, k)$ bidders with ironed virtual value greater than ℓ (so $|S| > k - \ell$). We next explicitly compute the payments collected by \mathcal{A}^* on this input, using the standard payment formula for incentive-compatible mechanisms (see [19] or [11]). Let a_i and b_i denote the left and right endpoints, respectively, of the interval of values v that satisfy $\bar{\phi}_i(v) = z$. Since the distributions are pointwise ordered, the a_i 's and the b_i 's are nonincreasing in i . Let $q = (k - \ell)/|S|$ denote the winning probability of a bidder in S . Define $q' = (k - \ell + 1)/(|S| + 1)$ as the hypothetical winning probability of a bidder in T if it lowered its bid to the value $\bar{\phi}_i^{-1}(z)$. The expected payment of a bidder i in S is qa_i (i.e., a_i in the event that it wins). The expected payment of a bidder i in T (who wins with certainty) is $q'a_i + (1 - q')b_i$. To complete the proof, we argue that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least the expected revenue collected by \mathcal{A}^* from the bidders in S , and also at least that from the bidders in T .

Projecting onto a subset of bidders only decreases the value of the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ (see Lemma A.1 for the formal argument). First, project onto the k bidders of S with the highest a_i values. Consider charging each such bidder the price a_i . This is a monotone price vector. By our assumption on the input \mathbf{v} , all of these prices are at most the second-highest valuation in \mathbf{v} . By the definitions, $v_i \geq a_i$ for every bidder $i \in S$ so every offer will be accepted. The resulting revenue is at least the expected revenue earned by \mathcal{A}^* on \mathbf{v} , and the value of the monotone price

benchmark can only be higher. This shows that $\mathcal{M}^{(2,k)}(\mathbf{v})$ is at least the expected revenue collected by \mathcal{A}^* from bidders in S .

Similarly, project onto the (at most k) bidders of T , and consider charging each such bidder i the price $q'a_i + (1 - q')b_i$. Again, this is a monotone price vector with all prices bounded above by the second-highest valuation of \mathbf{v} , and every offer will be accepted. The value of the monotone price benchmark can only be larger, so $\mathcal{M}^{(2,k)}(\mathbf{v})$ is also at least the expected revenue collected by \mathcal{A}^* from bidders in T . The proof is complete. ■

4.2 Reduction from Limited to Unlimited Supply

Having justified the k -unit monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$, we turn to designing auctions that approximate it well. We show that competing with this benchmark reduces to competing with the benchmark $\mathcal{M}^{(2)}$ in unlimited-supply settings. This reduction from limited to unlimited supply is a generalization of one for identical bidders [9]. The idea is to first identify the k “most valuable” bidders, and then run an unlimited-supply auction on them. In contrast to the identical-bidder setting in [9], the most valuable bidders with an ordering are not necessarily those with the highest valuations. For example, a high-valuation bidder late in the ordering need not be valuable, because extracting high revenue from it might necessitate excluding many moderate-valuation bidders earlier in the ordering.

We analyze the “black-box reduction” shown in Figure 6.

Input: A valuation profile \mathbf{v} for a totally ordered set $N = \{1, 2, \dots, n\}$ of bidders and k identical items. A truthful digital goods (unlimited supply) auction \mathcal{A} for ordered bidders.

1. Let p^* achieve the optimum monotone price benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ for \mathbf{v} and k . Let $S = \{i \in N : v_i \geq p_i^*\}$ be the set of winners under p^* .
2. Run the unlimited supply auction \mathcal{A} on the bidders S , with the induced bidder ordering.
3. Charge suitable prices so that truthful reporting is a dominant strategy for every bidder.

Figure 6: The auction BLACK-BOX REDUCTION (BBR).

Theorem 4.2 *If \mathcal{A} is a truthful unlimited-supply auction with ordered bidders that is β -competitive with $\mathcal{M}^{(2)}$, then the BLACK-BOX REDUCTION (BBR) auction is a truthful limited-supply auction with ordered bidders that is 2β -competitive with $\mathcal{M}^{(2,k)}(\mathbf{v})$.*

Proof: The BLACK-BOX REDUCTION (BBR) auction is clearly feasible, in that there are always at most k winners. The first step can be implemented efficiently using dynamic programming, so if \mathcal{A} runs in polynomial time, then so does the BLACK-BOX REDUCTION (BBR) auction. To see that it is truthful, first note that the allocation rule in the first step is monotone — if bidder i belongs to the computed set S in the profile \mathbf{v} , then it also belongs to S for every bigger valuation v_i (holding other bidders' valuations fixed). The composition of this rule with the truthful auction in the second step is also monotone (i.e., bidding higher only increases winning probability), so there are unique payments that render the auction truthful (see e.g. [11]). These payments are easy to describe. Monotonicity of the first step implies that each bidder i faces a threshold bid $t_i(\mathbf{v}_{-i})$ that is necessary and sufficient to be included in the computed set S . The payment of a winner i in the final step of the BLACK-BOX REDUCTION (BBR) auction is simply $t_i(\mathbf{v}_{-i})$ or the payment computed by \mathcal{A} , whichever is larger.

We prove the performance guarantee by arguing the following two statements: (i) the unlimited supply benchmark $\mathcal{M}^{(2)}$ applied to S is at least half of the limited-supply benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$ applied to the original bidder set; and (ii) the revenue of BLACK-BOX REDUCTION (BBR) on the original bidder set is at least that of the unlimited-supply auction \mathcal{A} with the bidders S . The second statement follows immediately from the facts that the winners of BLACK-BOX REDUCTION (BBR) are the same as those of \mathcal{A} , and that the winners' payments are only higher. For statement (i), consider prices p^* that determine the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. The projection p_S^* of this price vector onto the set S of bidders has revenue exactly $\mathcal{M}^{(2,k)}(\mathbf{v})$. If p_S^* is feasible, then it certifies that the benchmark $\mathcal{M}^{(2)}$ is at least $\mathcal{M}^{(2,k)}(\mathbf{v})$. The only issue is if the second-highest bidder is excluded from S , in which case p_S^* might use a price larger than the second-highest valuation in S (which is not permitted by the benchmark $\mathcal{M}^{(2)}$). But such a price can only extract revenue from the bidder with the highest valuation, and every price of p^* is at most the second-highest valuation $v^{(2)}$ of the original bidders. Thus, we can restore feasibility to p_S^* by lowering at most one price to the second-highest valuation of S , and we lose revenue at most $v^{(2)}$. Since $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq 2v^{(2)}$ — consider the price vector that offers $v^{(2)}$ to everybody — we retain at least half the revenue of p_S^* . Statement (i) and the theorem follow. ■

Of course, we can use the OPTIMAL PRICE SCALING (OPS) auction from Section 3 in Theorem 4.2 to obtain a truthful limited-supply auction that is $O(1)$ -competitive with the benchmark $\mathcal{M}^{(2,k)}(\mathbf{v})$. Theorem 4.1 implies that the resulting auction also enjoys a strong simultaneous approximation guarantee in Bayesian environments.

References

- [1] G. Aggarwal and J. D. Hartline. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1083–1092, 2006.
- [2] S. Alaei, A. Malekian, and A. Srinivasan. On random sampling auctions for digital goods. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 187–196, 2009.
- [3] P. Azar, C. Daskalakis, S. Micali, and S. M. Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.
- [4] J. Chen and S. Micali. Mechanism design with set-theoretic beliefs. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 87–96, 2011.
- [5] N. Devanur and J. D. Hartline. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 41–50, 2009.
- [6] N. Devanur, J. D. Hartline, A. R. Karlin, and T. Nguyen. A prior-independent mechanism for profit maximization in unit-demand combinatorial auctions. In *Proceedings of 7th Workshop on Internet & Network Economics*, 2011. To appear.
- [7] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, pages 129–138, 2010.
- [8] A. Goldberg and J. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, pages 29–35, 2003.

- [9] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [10] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999.
- [11] J. D. Hartline. Approximation in economic design. Book in preparation, 2012.
- [12] J. D. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.
- [13] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 75–84, 2008. Full version last revised August, 2012.
- [14] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 225–234, 2009.
- [15] J. D. Hartline and Q. Yan. Envy, truth, and optimality. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 243–252, 2011.
- [16] T. Ichiba and K. Iwama. Averaging techniques for competitive auctions. In *Proceedings of the Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pages 74–81, 2010.
- [17] S. Leonardi and T. Roughgarden. Prior-free auctions with ordered bidders. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 427–434, 2012.
- [18] G. Lopomo, L. Rigotti, and C. Shannon. Uncertainty in mechanism design. Submitted, 2009.
- [19] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

- [20] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9, pages 209–241. Cambridge University Press, 2007.
- [21] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 844–861, 2012.

A Missing Proofs

Lemma A.1 *For every valuation profile \mathbf{v} , $k \geq 2$, and subset S of the bidders with induced profile \mathbf{v}^S , $\mathcal{M}^{(2,k)}(\mathbf{v}) \geq \mathcal{M}^{(2)}(k, \mathbf{v}^S)$.*

Proof: (Sketch.) Fix an input \mathbf{v} , with monotone prices p^* determining $\mathcal{M}^{(2,k)}(\mathbf{v})$. By induction, we only need to show that adding a single new bidder i can only increase the value of the benchmark. Start by offering i the same price q as its predecessor in the ordering (or the second-highest valuation, if there is no predecessor). If i rejects (i.e., $v_i < q$), this extended price vector is feasible and we are done (the optimal feasible price vector is only better). If i accepts (i.e., $v_i \geq q$) then the price vector is infeasible (with $k + 1$ winners) and we argue as follows. Go through the bidders after i one by one, increasing the offer price to q . This preserves monotonicity. If a previously winning bidder ever rejects this higher offer price, we are done (feasibility is restored and the overall revenue is higher). If not, there is now a “suffix” of bidders with the common offer price q . (This case only occurs if i is after all of the winners in p^* .) We now increase their common offer price until it equals that of the previous bidder, thereby increasing the number of bidders in the suffix. Eventually a bidder that was winning under p^* will reject the new offer price (otherwise it would contradict the optimality of p^*), leaving us with a feasible monotone price vector with revenue at least that of the original one. ■